# SciDAC's Earth System Grid Center for Enabling Technologies

***Semi-Annual Progress Report for the Period***
***April 1, 2010 through September 30, 2010***

***Principal Investigators***
Dean N. Williams[3], Ian T. Foster[1], and Don E. Middleton[6]

***The Earth System Grid Center for Enabling Technologies Team:***
Rachana Ananthakrishnan[1], Neill Miller[1], Frank Siebenlist[1]
Arie Shoshani[2], Alex Sim[2],
Gavin Bell[3], Robert Drach[3], Michael Ganzberger[3],
Jim Ahrens[4], Phil Jones[4],
Daniel Crichton[5], Luca Cinquini[5],
David Brown[6], Danielle Harper[6], Nathan Hook[6], Eric Nienhouse[6], Gary Strand[6], Hannah Wilcox[6],
Nathaniel Wilhelmi[6], Stephan Zednik[6],
Steve Hankin[7], Roland Schweitzer[7],
Meili Chen[8], Ross Miller[8], Galen Shipman[8], Feiyi, Wang[8],
Peter Fox[9], Patrick West[9], Stephan Zednik[9],
Ann Chervenak[10], and Craig Ward[10]

**Climate simulation data are now securely accessed, monitored, cataloged, transported, and distributed to the national and international climate community**

[1] Argonne National Laboratory
[2] Lawrence Berkeley National Laboratory
[3] Lawrence Livermore National Laboratory
[4] Los Alamos National Laboratory
[5] National Aeronautics and Space Administration
[6] National Center for Atmospheric Research
[7] National Oceanic and Atmospheric Administration
[8] Oak Ridge National Laboratory
[9] Rensselaer Polytechnic Institute
[10] University of Southern California, Information Sciences Institute

# Contents

# 1 Executive Summary

This report summarizes work carried out by the Earth System Grid Center for Enabling Technologies (ESG-CET) during the period April 1, 2010 through September 30, 2010. It discusses highlights, overall progress, period goals, and collaborations and lists papers and presentations. To learn more about our project and to find previous reports, please visit the ESG-CET websites (URL: http://esg-pcmdi.llnl.gov/, https://wiki.ucar.edu/display/esgcet/Home). This report will be forwarded to program managements for the Department of Energy (DOE) Scientific Discovery through Advanced Computing (SciDAC) and the Office of Biological and Environmental Research (OBER), to national and international collaborators and stakeholders (e.g., the Coupled Model Intercomparison Project, phase 5 [CMIP5] for the Intergovernmental Panel on Climate Change [IPCC] 5th Assessment Report [AR5]; the Community Climate System Model [CCSM]; the Climate Science Computational End Station [CCES]; SciDAC II: A Scalable and Extensible Earth System Model for Climate Change Science; the North American Regional Climate Change Assessment Program [NARCCAP]), and to researchers working on a wide-range of other climate model and observation evaluation activities.

The ESG-CET executive committee consists of Dean N. Williams, Lawrence Livermore National Laboratory (LLNL); Ian Foster, Argonne National Laboratory (ANL); and Don Middleton, National Center for Atmospheric Research (NCAR). The ESG-CET team is a group of researchers and scientists with diverse domain knowledge, whose home institutions include eight laboratories and two universities: ANL, Los Alamos National Laboratory (LANL), Lawrence Berkeley National Laboratory (LBNL), LLNL, National Aeronautics and Space Administration (NASA)/Jet Propulsion Laboratory (JPL), NCAR, Oak Ridge National Laboratory (ORNL), Pacific Marine Environmental Laboratory (PMEL)/National Oceanic and Atmospheric Administration (NOAA), Rensselaer Polytechnic Institute (RPI), and University of Southern California, Information Sciences Institute (USC/ISI). All work is accomplished under the DOE open-source guidelines and in close collaboration with the project's stakeholders, domain researchers, and scientists.

The Earth System Grid (ESG) project has developed and delivered a production environment that serves a worldwide community with climate data from multiple climate model sources (e.g., CMIP [IPCC], CCSM), ocean model data (e.g., Parallel Ocean Program), and analysis and visualization tools. Data holdings are distributed across multiple sites including LANL, LBNL, LLNL, NASA, NCAR, and ORNL as well as unfunded partners such as Geophysical Fluid Dynamics Laboratory/NOAA. At LLNL, ESG operates a dedicated portal that supports the IPCC community in developing the Coupled Model Intercomparison Project (CMIP). We estimate that over 700 scientific publications focus on analyses of IPCC data alone. By the end of 2010, the combined number of registered users from all ESG portals is likely to exceed 25,000. ESG also crossed a milestone this year by delivering more than *1 petabyte* (PB) of data to its users and is currently approaching 1.2 PB of downloaded data.

More recently, ESG has focused on broadening and generalizing the ESG system to support a more distributed, international, and diverse collection of archived sites and data types. In addition, ESG is extending services beyond access to raw data files by developing information products (scientific graphics, animations, etc.) and "server-side analysis" capabilities that will allow users to request output from commonly used analysis and intercomparison procedures.

# 2 Reporting Period Highlights

Our overall goals for this reporting period were to prepare for receiving and disseminating the CMIP5 (IPCC AR5) data and to move forward with other project missions, such as product services and client analysis tool access.

The ESG-CET is intended to serve customers who span a broad spectrum of sophistication. These users range from numerical modelers who want access to "raw" model output files and verbatim subsets of model output, to climate impacts investigators who want rapid access to these data without the complexities of model-specific coordinate systems, and users who only want to quickly visualize the overall behaviors of models. The petascale nature of the ESG data holdings requires that significant levels of data reduction take place at the server so we can satisfy these customers—both through straightforward subsetting and decimation and through specific analysis operations, such as the computing of spatio-temporal averages. In the ESG-CET architecture, we refer to the steps that convert raw data into analysis results and visualizations as "product services."

## 2.1 Highlights

This section describes the team's accomplishments during the reporting period.

### *2.1.1 ESG Release Schedule*

Since the last reporting period, the ESG team has moved beyond beta releases and release candidates to official release versions 1.0 and 1.1. Our next scheduled release (version 1.2) is targeted for October 20, 2010. This baseline CMIP5 release will contain features for user interface support, tokenless authorization, full Data Reference Syntax (DRS) metadata support, and basic CMIP5 XML harvesting from the questionnaire developed by the European Union's (EU's) METAFOR (Common Metadata for Climate Modelling Digital Repositories) project, to name a few. **Table 1** provides more details on this and future releases. In reviewing the table, note that release schedules serve as a guideline, not as a strict plan for releasing the software, because the updates released by the ESG team consist of the disparate complex software packages that make up the Gateway and the Data Node.

**Table 1**: Source code release schedule of the ESG enterprise distributed system.

| **1.2.0-CMIP5 Baseline 1** | **Goal:** Minimum requirements needed for CMIP5 data archiving and dissemination | **Date:** Monday, October 20, 2010 |
|---|---|---|

**Gateway**

- Basic replication.
  - UI support for replicated dataset discovery and script generation.
  - Harvesting of replication-related Thredds catalog properties (master_gateway, is_replica).
- X509 certificate and OpenID based authorization (AKA: Token-less authorization.)
  - Enable OpenID and certificate based auth-z in datanode services (TDS, GridFTP/DML.)
  - Enable gateway attribute service.
  - WGet script generation w/ x509 certificates.
  - Configurable DataNode services based to enable/disable token generation.
  - MyProxyLogon WebStart integration.
- Full DRS metadata ingestion from Thredds Catalogs.
  - Includes new field cmor_table
- Automatic OAI RDF harvesting.
  - Administrative UI for schedule configuration.
  - Integrate manual harvesting into scheduled process.
- Basic CMIP5 XML harvesting from METAFOR questionnaire atom feed (may be bumped to 1.3)
  - May include QC related fields (TBD)
  - **Note:** Dependent on stable/finalized atom feed.
- Streamline Gateway installation and update.
  - Simplified application configuration files.
  - Improved documentation.
- Gateway CA and Host certificate upgrade (Federation-wide deployment issue.)
- Whitelist support for federation services
  - Auth-z services and consumers
  - Attribute services and consumers

**Data Node**

- Simplified Data Node Installation Process
- New Security using SAML and OpenID (tokenless)
- Secure Attribute services support using SAML with Yaris service
- DRS support for file system layout
- Solr search support

| **1.3.0-CMIP5 Baseline 2** | **Goal:** Enhance ESG user experience by providing minimum product services | **Date:** Monday, January 17, 2011 |
|---|---|---|

**Gateway**

- Security Enhancements (Enumerated in separate document)
- Enhance data download workflows:
  - Improved script generation (via Data Cart capabilities, add wget help, version reqs.)
  - Improved large fileset selection
- LAS Product Server integration (datasets/variable selection with auth-z.)
  - Enable basic sub-setting services.
  - Enable basic data visualization services.
- Manage file access locations using publisher:
  - Support for file rename, move and remove use cases.
- Gateway collation of datanode manager metrics:
  - Define Service API for gatewaydatanode metrics transfer.
  - Implement automated process (and UI) for routine datanode query.
  - Ensure datanode file access services persist download metrics

    **Note:** Dependent on full file download metrics capture and persistence at datanode including GridFTP (DN)

**Data Node**

- Catalog canonicalization via the DRS
- DRS / Filesystem indirection layer for support of novel file systems
- Virtual Machine Packaging of Data Node (OVF descriptor to stand up Centos JeOS VM)
- Enhanced Per Node Metrics Interface
- P2P node discovery and ad-hoc overlay mesh network of nodes
- Distributed index search across domains.

| 1.5.0-CMIP5 – Baseline 3 | **Goal:** Support for CMIP5 distributed archive and workflow | **Date:** Monday, March 8, 2010 through Monday, April 19, 2010 |
|---|---|---|

Gateway

- Data access by DRS URLs
- Advanced support for replica discovery and download
  - Search of replicated datasets
  - Wget script generation at remote Gateway
- Present dataset version history in UI.
- Enhanced CMIP5 XML harvesting from METAFOR questionnaire atom feed.
- Fully support deletion of datasets
  - Verify database integrity after hard deletion
  - Support logical deletion
- Optional token support discontinued based on federation feedback.

Data Node

- Publisher graphics user interface

### *2.1.2 Development and Operation of ESG Infrastructure for Observational Data Sets*

ESG team members from NASA, LLNL, ORNL, and NOAA are collaborating with the goal of incorporating observational data sets into the CMIP5 global archive and services so that they can be used as part of the upcoming IPCC report. The following progress is worth mentioning:

- A new NASA/JPL system composed of an ESG Gateway and Data Node has been deployed and is now operational at JPL. This system was federated with the other ESG Gateways through the exchange of certificate trust roots and cross-site linking.
- Selected NASA satellite data collections (AIRS, MLS, TES) were converted to the Network Common Data Format - Climate and Forecasting (NetCDF-CF) conventions and published to the system. Work on including other collections is under way.
- All parties interested in publishing observations have started to collaborate on defining a common strategy for formatting, naming, organizing, and documenting the observational data sets to maximize their usefulness and value to the scientific community. We are experimenting with running CMOR2 for compliance tests with the model output, and we are writing a best-practices document to guide other agencies in contributing observations to the CMIP5 archive.
- JPL in collaboration with other NASA groups has started work on defining a long-term architecture to integrate the ESG infrastructure with existing data services operated by the NASA Distributed Active Archive Centers (DAACs).

ORNL has established a pilot program to share observations to support model-to-data comparisons. The goal is to demonstrate the feasibility of integrating high-value data sets within ESG such as Atmospheric Radiation Measurement (ARM), Carbon Dioxide Information and Analysis Center (CDIAC), and ORNL DAAC. These data sets will provide access to a wealth of DOE observations via ESG. Prototypes to exploring such integration have initially produced promising results. However, integration of these data sets will require substantial efforts in spatio-temporal projections, metadata harvesting and integration, and observational and model data comparison tools. (See section 2.1.2, “ESG Challenges Requiring Further R&D Efforts,” for more detailed information.)

### 2.1.3 Data Download Highlights

Figure 1 shows the ESG download metrics since 2005, which include data download support for CMIP3, CCSM, and NARCCAP as well as the Parallel Ocean Program (POP), Parallel Climate Model (PCM), and Carbon Land Model Intercomparison Project (C-LAMP). Figure 1 also shows, ESG-CET passed the petabyte mark for total data delivered early in 2010.

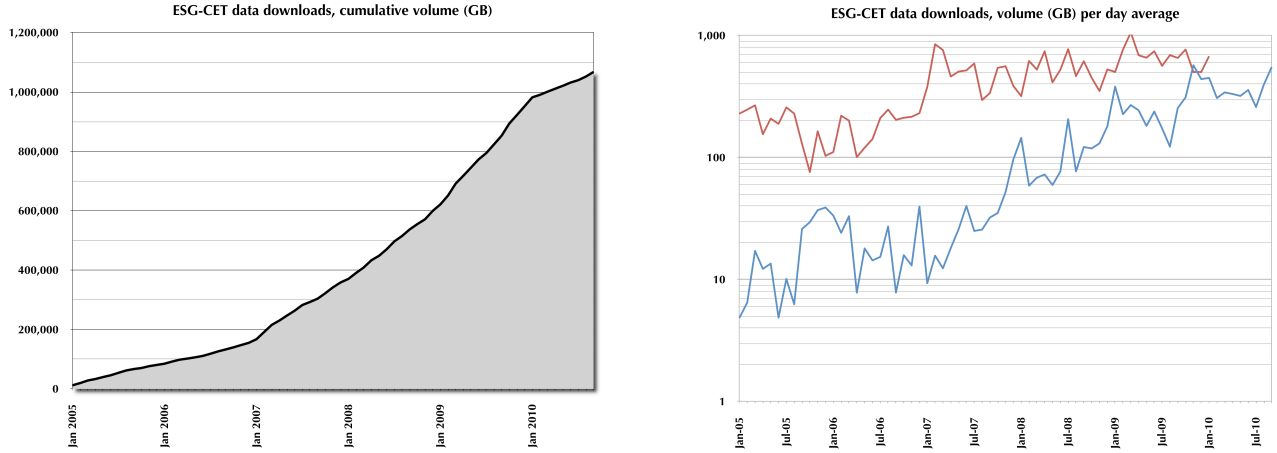


**Figure 1:** ESG accumulated downloads, well over 1 PB.

### 2.1.4 LLNL Prepares for CMIP5 (IPCC AR5) Data

In preparation for the CMIP5 data, the ESG team has reached out to multiple CMIP5 climate model centers. Of the sites, the Institut Pierre Simon Laplace des sciences de l'environnement (IPSL) and the Centre National des Recherches Météorologiques (CNRM) in France, the NASA Center for Climate Simulation (NCCS) and the Geophysical Fluid Dynamics Laboratory (GFDL) in the U.S., and the Centro Euro-Mediterraneo Cambiamenti Climatici (CMCC) in Italy have set up ESG Data Nodes and published test model data to the LLNL Gateway. In addition, other non-CMIP5 sites have started efforts to install ESG Data Nodes. CMIP5 data federation and dissemination will begin in October.

In a truly collaborative effort, the ESG team has worked with the Deutsches Klimarechenzentrum (DKRZ) and the British Atmospheric Data Centre (BADC) to flesh out access to the CMIP5 data, including requirements for Data Reference Syntax (DRS), Quality Control (QC), Digital Object Identifiers (DOIs), and modeling group questionnaire input. Figure 2 shows the data flow before and after the questionnaire is available. After data are published the ESG publisher automatically performs the QC level-1 (QCL1) check and flags the published set as QCL1-compliant. Next, the DKRZ QC code is run to check for QC level-2 (QCl2) compliance, and the ESG system flags the data as such. Finally, a visual inspection is performed to check for inconsistencies and metadata correctness. If the data are in compliance with QC level-3 (QCl3), they are flagged accordingly, and a DOI is issued.

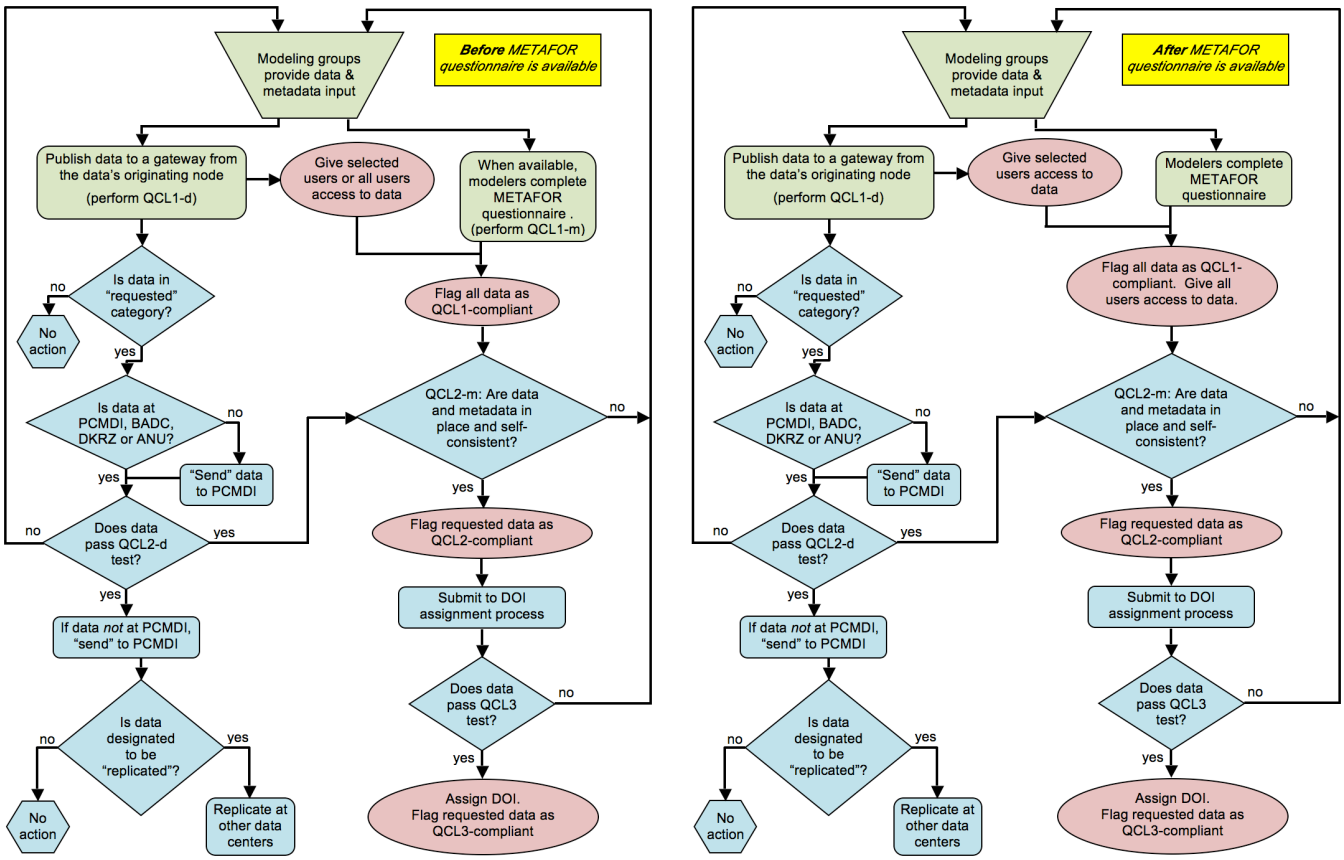


**Figure 2:** CMIP5 data flow diagram, courtesy of Karl Taylor (LLNL), Bryan Lawrence (BADC), and Michael Lautenschlager (DKRZ).

### *2.1.5 NASA ESG Gateway R&D Highlights*

The ESG team, specifically at JPL, has deployed and started operating a new Gateway and Data Node system dedicated to serving NASA observational data sets to the CMIP5 (IPCC AR5) community. A few selected satellite data sets (AIRS, TES, MLS—shown in Figure 3) have been published to the ESG system. ANL, LLNL, JPL, and BADC have worked to define and adopt common security protocols and application programming interfaces (APIs) as the basis for interoperability within the global ESG Federation, and to guarantee seamless user access to the CMIP5 data archive. Java and Python implementations of the ESG security protocols have been developed, deployed and tested. The ESG NASA Gateway hosted at JPL is available at: http://esg-gateway.jpl.nasa.gov/home.htm.

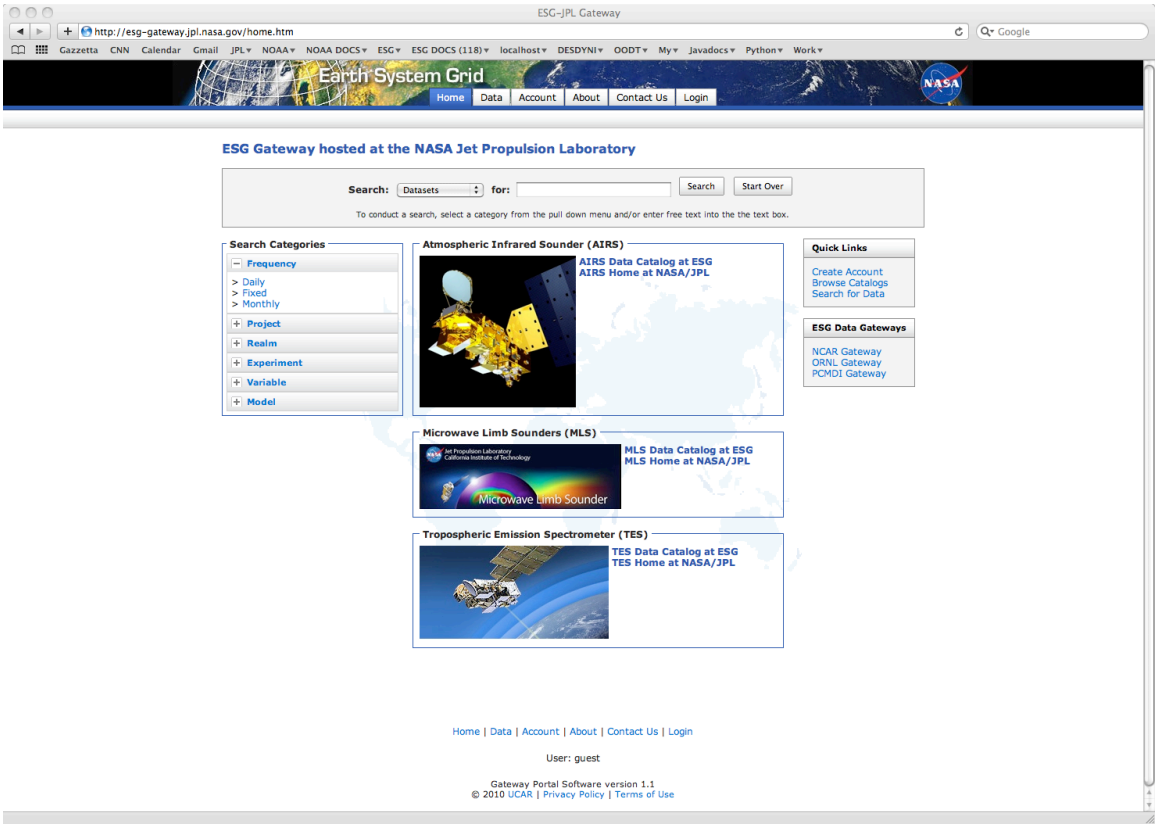


**Figure 3:** Home page for the new NASA JPL Gateway.

### *2.1.6 NCAR ESG Gateway Portal R&D Highlights*

ESG accomplished several critical milestones for its Gateway during this reporting period, including two major releases along with important improvements in software development management and community engagement. We released the ESG Gateway Version 1.0 in May and executed the full production portal cutover at NCAR on May 26, 2010. This process involved migrating more than 16,400 user accounts and 1.3 million files to the new system. It went remarkably smoothly, with very few user problems. In parallel with this production cutover at NCAR, we released the new system and deployed it for evaluation and assessment purposes at LLNL, ORNL, BADC, and DKRZ. Version 1.0.0 provided many additional features, improvements, and fixes including versioned data sets, access-controlled file downloads, deep-storage transfers, data publishing performance improvements, and user interface review and refinements.

We proceeded to release V1.1.0 and V1.1.1, which offered more important features for ESG (see section 3.1 *"Gateway Software Development"* for greater details), and then focused on the development work required to release the critical CMIP5 Baseline release, V1.2.0, in early fiscal year (FY) 2011. Since the production ESG cutover at NCAR in May 2010, we have added over 1,500 new users to the ESG system and published over 400,000 new files. The NCAR ESG Gateway now has 1.7 million published files, representing more than 400 terabytes (TB) of available climate data.

In addition to the development and deployment work, we significantly improved the software engineering infrastructure and process and took major steps toward establishing a community development environment. Also noteworthy is that all of this work reflects extensive collaboration among the core ESG-funded partners, the NOAA Global Interoperability Program (GIP) and Curator effort; the EU METAFOR project, and the international Global Organization for Earth System Science Portals (GO-ESSP), especially BADC and DKRZ.

### *2.1.7 ORNL ESG Gateway Portal R&D Highlights*

The ORNL ESG-CET team successfully deployed the next-generation ESG infrastructure to support the needs of the climate-modeling community using the world's most powerful simulation environment at the Oak Ridge Leadership Computing Facility (OLCF) and users of this data from around the world. ORNL helped develop, test, and deploy the ESG system in collaboration with the entire ESG-CET team. Notable highlights include transitioning C-LAMP from our legacy infrastructure to the next-generation ESG infrastructure. The ORNL team also developed sophisticated data publication mechanisms allowing full metadata captures for data sets in deep-archive systems, thus facilitating the design and development of data-access components for sites with multiple security enclaves. These features will be critical to providing the climate science community with access to climate modeling data sets generated at OLCF from CCSM activities.

### *2.1.8 LANL ESG Node Highlights*

LANL has worked to have the latest ESG Data Node software installed on the new Oceans11 hardware and is currently working to republish the POP data onto the new system. Once work is completed, LANL will start adding new ocean data to the ESG archive. Because of the Direct Attached Storage, we anticipate the new Ocean11 machine will fix the download problems experienced in large multiple data transfers.

In the last progress report, we mentioned the beginning development of the netCDFPOP reader and ParaView visualization efforts. The netCDFPOP reader is now in beta release, and we are working with a prototype ParaView POP visualization tool. Figure 4 shows an example of netCDFPOP temperature ocean output visualized with ParaView. In addition, a beta version of the Spherical Coordinate Remapping and Interpolation Package (SCRIP) software is now available for download. This software is more robust than the previous interpolation package and includes thread parallelism to handle larger grids. We are also adding a normalization feature to accommodate use of model-computed areas (rather than SCRIP-computed grid areas). Once completed, we will officially release SCRIP to the user community.

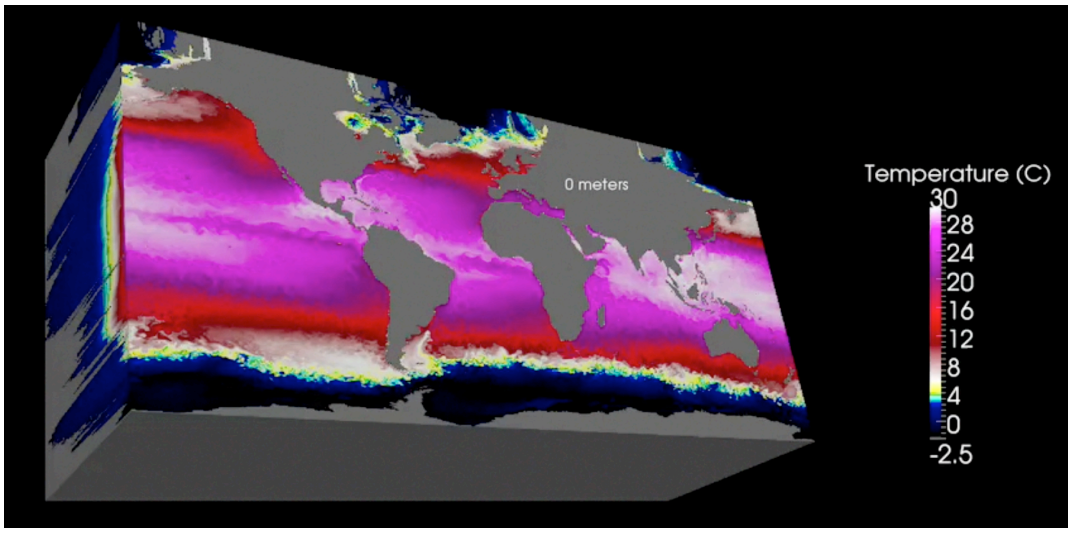


**Figure 4:** Read from the ESG POP archive, this image is the first in a sequence using a ParaView server running at ORNL and ParaView client running at LANL.

### *2.1.9 LBNL Berkeley Storage Manager (BeStMan)*

IPCC AR4 CMIP3 data sets (about 36 TB) have been replicated from LLNL to the National Energy Research Scientific Computing Center (NERSC) using the Bulk Data Mover (BDM). Figures 5 and 6 show the transfer throughput performance (in seconds) for BDM data transfers. The end-to-end throughput performance was measured while transferring 17,049 files in 5.3-TB and 16,604 files in 5-TB of climate data sets with 30 concurrent transfers. Transfer throughput was consistent most of the time throughout the request, as expected, showing effectiveness of the transfer management in BDM. However, as shown in Figure 6, we detected low performance in the middle of the transfers, while the number of concurrency was still at 30 altogether. This issue caused each concurrent connection performance to be much lower and may have caused packet loss, too. The adaptive transfer management, now being added to BDM, will help resolve this problem by minimizing overhead of slow data transfers during the low-performance period. When BDM identifies a performance slowdown, it will reduce the number of concurrent transfers dynamically to maximize the per-connection throughput and thus maximize the resource usability during the transfer. This feature will be available in a future release of BDM.

The NERSC Gateway and Data Node have been deployed, serving NERSC climate data sets and CMIP3 data sets.

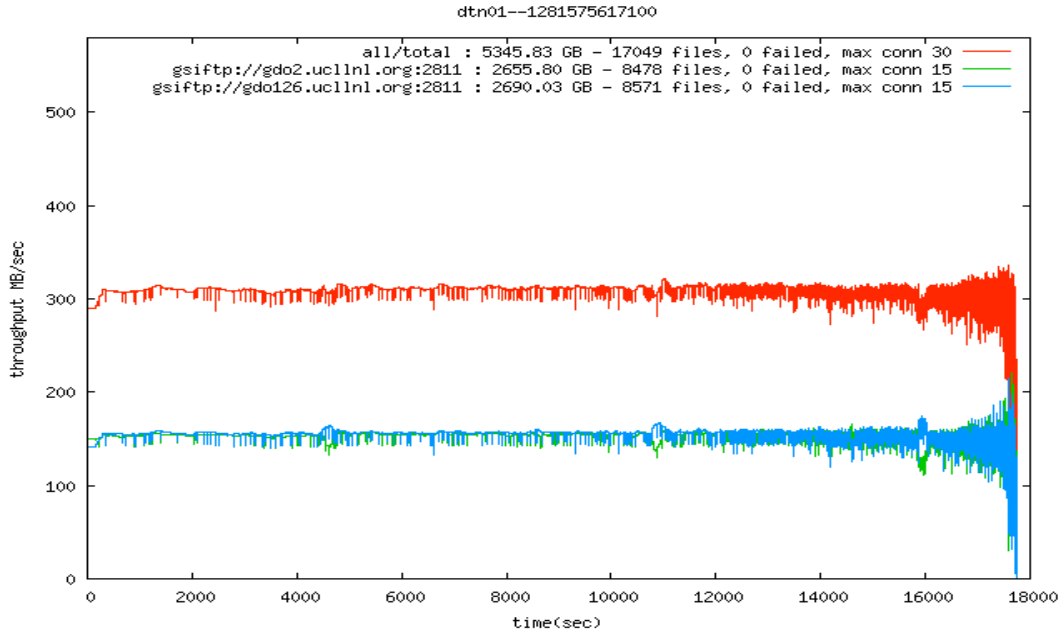


**Figure 5:** Results for a BDM transfer show transfer throughput over time in seconds. This operation involved transferring 17,049 files in 5.3-TB of climate data sets from two sources at LLNL to one destination at NERSC with 15 concurrency for each data source on a shared network.

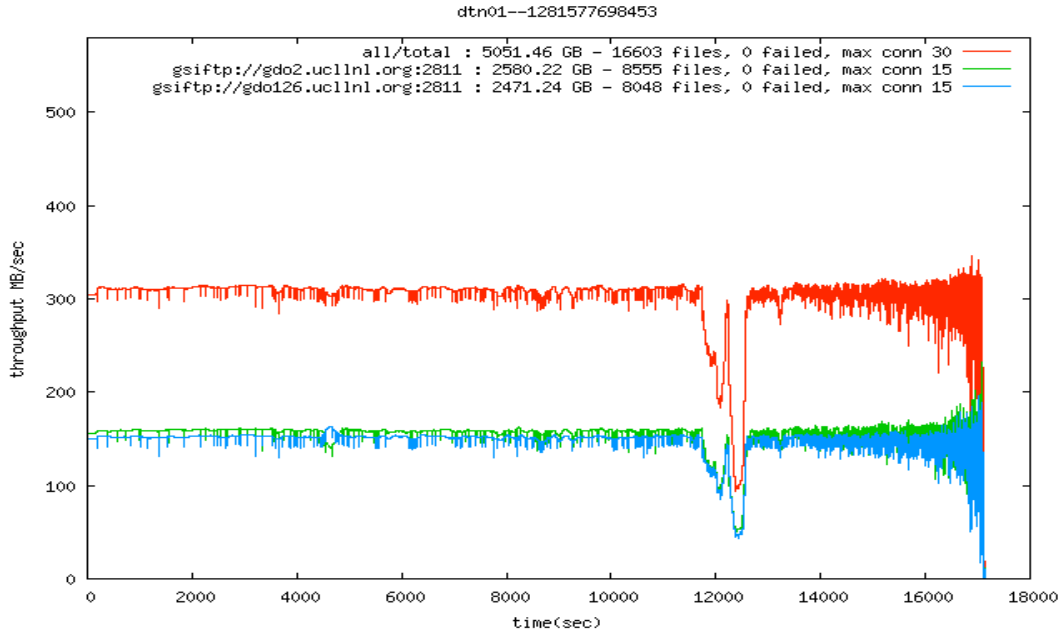


**Figure 6:** Results for another BDM data transfer, this one from two source servers at LLNL to one destination at NERSC for 16,603 files in 5-TB of CMIP-3 data sets.

### 2.1.10 PMEL Product Delivery Services Highlights

At the heart of PMEL's contribution to the ESG-CET lies LAS, a development activity with origins that date back to 1993. LAS is an XML-configurable workflow engine that supports widely varied climate applications. The incorporation of LAS into ESG in 2006 brought added workflow support to the ESG project and expanded the positive impacts of ESG-specific enhancements to LAS. Since the previous reporting period, the PMEL team has enhanced the model-intercomparison capabilities of LAS. In addition to the ongoing work on the model inter-capabilities (which have been enhanced to work with vectors), PMEL has completed several significant milestones for this reporting period. Additionally, the LAS team has worked with collaborators at NASA and ORNL to finish the first phase of an effort to integrate the LAS servers and the LAS JavaScript browser client with the Authentication and Authorization infrastructure used by the ESG-CET project.

### 2.1.11 ANL Security, Data, and Services Highlights

Argonne proposed and led an effort to define federation-wide security acceptance tests that articulate the compliance requirements for all deployments on the federation. Follow-up efforts are planned to implement and package these tests for verification and validation of the federation infrastructure. These tests will be integrated into an automated monitoring service that reports on the compliance and validity of the services in the federation.

Argonne is also working with the ESG team on multiple efforts to test and patch various pieces of the software, to prepare for production release. Specifically, we deployed and tested new data node components for security THREDDS Data Server (TDS) and clients for GridFTP and BDM. We are also performing integrated testing of the different software pieces to improve quality and assurances on the deployed software.

In addition, ANL is preparing the security infrastructure required for accepting and distributing CMIP5 data and enabling seamless access to users across the ESG federation. These efforts were focused on two key areas: data movement using GridFTP servers and federation security.

***GridFTP: High-Performance Data Movement:*** Efforts to provide GridFTP as an option for end-user download resulted in integrating mechanisms to support GridFTP deployments on data node with ESG security plug-ins. This feature allows any user who has registered with the ESG infrastructure to access the GridFTP server to download data. Authorization plug-ins call out to the ESG authorization service, determine if a registered user has access to the particular data requested. In addition, we worked with the Globus GridFTP team to provide a restricted data root functionality to the GridFTP installation. This feature, which has been incorporated with the ESG Data Node installation script, allows data providers to choose the virtual root space to which GridFTP has access and thus restricts the data an end user can access on the machine.

Another important aspect of the ESG infrastructure is collecting usage information to report data download metrics. Argonne designed and developed a metrics reporting plug-in that records information about the data downloaded from a GridFTP server. The architecture leverages existing options and cleanly layers onto the GridFTP options. The plug-in uses the Globus Toolkit–enabled usage-reporting mechanism to capture data needed by ESG and transfer them to metrics databases on the node. We worked with the others on the team to integrate the metrics collection options with the installation scripts, such that they are deployed with any data node install.

***ESG Federation:*** The trust model and federation policy are critical in defining the federation and binds the independent sites as a single federation. Argonne led the efforts, with significant contributions from

collaborators at BADC, to document a Security Interface Control Document that defines the agreed contract for security aspects of the federation. This document explains the various security services and the agreements on VO attributes and namespaces for the use by the federation.

To facilitate distribution and coordination of trusted roots in the federation, Argonne set up a central repository of federation trust roots. The repository consists of trusted Certificate Authorities (CAs), Certificate Revocation Lists for the CAs, and Signing Policy for the CAs. In addition, the repository distinguishes the trusted CAs to be used for Public Key Infrastructure (PKI) authentication and the OpenID providers, and ESG nodes to provision the machines with the latest trust root information can use it. The trust root information is supported in two formats, Privacy Enhanced Mail (PEM) files and Java Keystores, for use by Globus tools and Java application servers. We are currently adding security information such as trusted attribute services, authorization services, and Gateways and Data Nodes for use in white list authorization. This information is provisioned in the ESG Gateway and can also be accessed by clients using the MyProxy Logon functionality.

### 2.1.12 USC/ISI Replication Highlights

The data replication team at ISI successfully implemented the ESG replication client during this period and completed end-to-end testing of this client, which mirrors data sets from LLNL to two sites: USC/ISI and NCAR. The replication client integrates functionality from several ESG providers. It queries the metadata catalog at the Gateway, pulls the THREDDS catalog from the original publication site's data node, and creates control files for either the BDM client or the GridFTP globus-url-copy client to transfer data. It also invokes the publication API from LLNL to scan the replicated files, generate a THREDDS catalog for those files, and publish the mirrored data set to the Gateway. The replication client compares the THREDDS catalog generated for the newly copied files with that created at the original publication site; the data set is only published if this comparison verifies the correctness of the THREDDS catalog. In tests of the completed replication client, it successfully provided end-to-end data replication of a 5- and 50-GB data set from LLNL to data nodes at USC/ISI and NCAR.

### 2.1.13 RPI Highlights

To participate fully in ESG development, RPI made progress in porting the ESG Data Node to RPI's new development environment for code advancement and testing—a dedicated Linux server supporting ESG software. This effort is needed for direct OPeNDAP Hyrax security, product services, GridFTP, and THREDDS catalog integration within ESG. In the process, RPI stood up a MyProxy server to authenticate RPI-based test users and now manages its own user access permissions for security reasons. Work is now under way to install an ESG Gateway at RPI in anticipation of developing the client–server relationship between client analysis tools that access ESG archives. Optimization work is needed to connect clients to the OPeNDAP Back-End Server (BES) and limit the number of connections. Design work to create a secure connection to BES for all transmissions, not just client authentication, will allow the passing of X.509 certificates to the BES on behalf of the client. It will also allow RPI to secure data transfer for sensitive data. RPI is also creating an administrative interface to the BES to improve connection maintenance, debugging work, starting and stopping the server, and dynamic configuration of the server.

We are installing and becoming acquainted with various systems that data users are using to visualize their work, such as TDS, FTDS, and THREDDS catalogs. This work will improve support for THREDDS catalogs in the OPeNDAP community by moving all functionality to the BES from its current location in the OPeNDAP Light-Weight Front-End Service (OLFS). The effort includes design

work on distributed data access and data manipulation using the BES, for example, distributing access and aggregation functionality. It also involves design work on a feature to register server-side functions for constraining, projecting, and selecting data, which will allow users to build these functions that can be dynamically loaded by the BES. To help move this effort forward, RPI is organizing and hosting a meeting with core OPeNDAP developers between October 6 and 8, 2010.

RPI will also continue its efforts with the Remote NetCDF Interface (RNI). Included in this work will be upgrading the RNI GridFTP module to work with the latest version of the BES communication layer. The RNI BES module will also be upgraded to work with the latest versions of the OPeNDAP libdap++ and the BES interface libraries.

# 3 Component Progress

During this reporting period, we made progress in the key areas necessary to meet the ESG-CET objectives, goals, and milestones. This section describes that work in greater technical depth and presents the components needed for the CMIP5 baseline release.

## 3.1 Gateway Software Development

ESG accomplished several major milestones for its Gateway, including two major releases. In addition, we improved the process for software development management and community engagement.

***ESG Gateway Version 1.0.0:*** We released the ESG Gateway Version 1.0 in May 2010 and executed the full production portal cutover at NCAR on May 26. This process involved migrating more than 16,400 user accounts and more than 1.3 million files to the new system. It went remarkably smoothly, with very few user problems, in part because our global user community participated in a high level of testing and evaluation prior to the transition. In parallel with this production cutover at NCAR, we released the new system and deployed it for evaluation and assessment at LLNL, ORNL, BADC, and DKRZ. Version 1.0.0 included many additional features, improvements, and fixes, such as versioned data sets, access-controlled file downloads and deep storage transfers, data publishing performance improvements, and user interface review and refinements.

***ESG Gateway Version 1.1.0:*** We proceeded to release Version 1.1.0, which addressed OpenID logins, data-set replication, Gateway Open Archives Initiative (OAI) Resource Description Framework– (RDF-) based metadata exchange, application build system upgrades, and a range of security improvements. This upgrade was followed by a minor release, Version 1.1.1.

***Progress towards ESG Gateway Version 1.2.0:*** As we progressed toward the critical Version 1.2.0 CMIP5 Baseline release early in FY 2011, we further addressed OpenID authentication support, completion of core replication functionality, and enhancements that support the ESG federation capabilities, including automated, scheduled OAI RDF metadata exchange and implementation of a federation service trust store and white-list support. We have also improved our data access authorization model, addressing X509 certificates and OpenID-based data access. We enhanced CMIP5 DRS support, providing for full DRS metadata harvesting by the publishing system and DRS-based metadata search and facet-based navigation. Collaborating extensively with NOAA GIP, Curator, and the EU-based METAFOR project, we continued to develop capabilities for the “model trackback” interface and the corresponding METAFOR Conceptual Information Model (CIM) ingest capabilities. On a closely related note, we are also working to address CMIP5 QC requirements.

In the 1.1.0 release, many site-specific configuration details remained as part of the core gateway component. As a result, the deploying sites and NCAR had to accommodate these customization and deployment properties. We executed a major overhaul of the software, centralizing configuration properties outside the gateway for easier maintenance. We also simplified the customization files and improved documentation for application installation and update. These improvements have been well received by the international community thus far. Lastly, we set up a new distribution format for the gateway.

***Gateway Software Engineering Infrastructure, Process, and Community Engagement:*** During the reporting period, we continued to improve our software engineering infrastructure and process so that the developer community can be more engaged in software development. We refined our use of the Atlassian system for issue-tracking and release control, addressing 201 specific issues over the past 6 months. We integrated Nexus as the repository manager for all software releases and instituted automated nightly builds of the gateway using the Atlassian Bamboo framework. We also refactored the Subversion repository to support community development and shared the repository with a small group of U.S. and international collaborators. Work on the Liquibase approach for cross-release database migration continues to be a foundation of enabling easy, rapid deployment at existing ESG Gateway sites.

We implemented practical and required coding standards for the Gateway software base and reduced the number of violations from over 8,000 to 1,700. The intent is to tighten the standards so we can continue to reduce the number. This effort will allow us to deploy a software system that the user community can maintain over the long term. Because the ESG Gateway has many dependencies, we switched to the Maven build system to improve dependency management while simplifying build functionality. We completed a Contributor License Agreement (CLA) and began executing it with developers from ORNL, BADC, and JPL/Caltech. CLAs provide a clean mechanism for enabling contributions to the Gateway codebase. Work continued to complete the process of having the entire Gateway under an OSI-approved Open Source License in the first quarter of FY 2011. Although the analysis is ongoing, we expect is that this license to be the Berkeley Software Distribution (BSD) derivative instead of the Apache model we originally targeted.

***Metrics and Summary:*** All of the work described above reflects extensive collaboration among core ESG-funded partners, the NOAA GIP effort, Curator, the EU METAFOR project, and the international GO-ESSP community, especially BADC and DKRZ. In summary, since the production ESG cutover at NCAR in May 2010, we have added over 1,500 new users to the ESG system and published over 400,000 new files. The NCAR ESG Gateway now has 1.7 million published files, representing over 400 TB of available climate data. The user community continues to ingest larger files on average, and usage continues to steadily climb (see Figure 7).

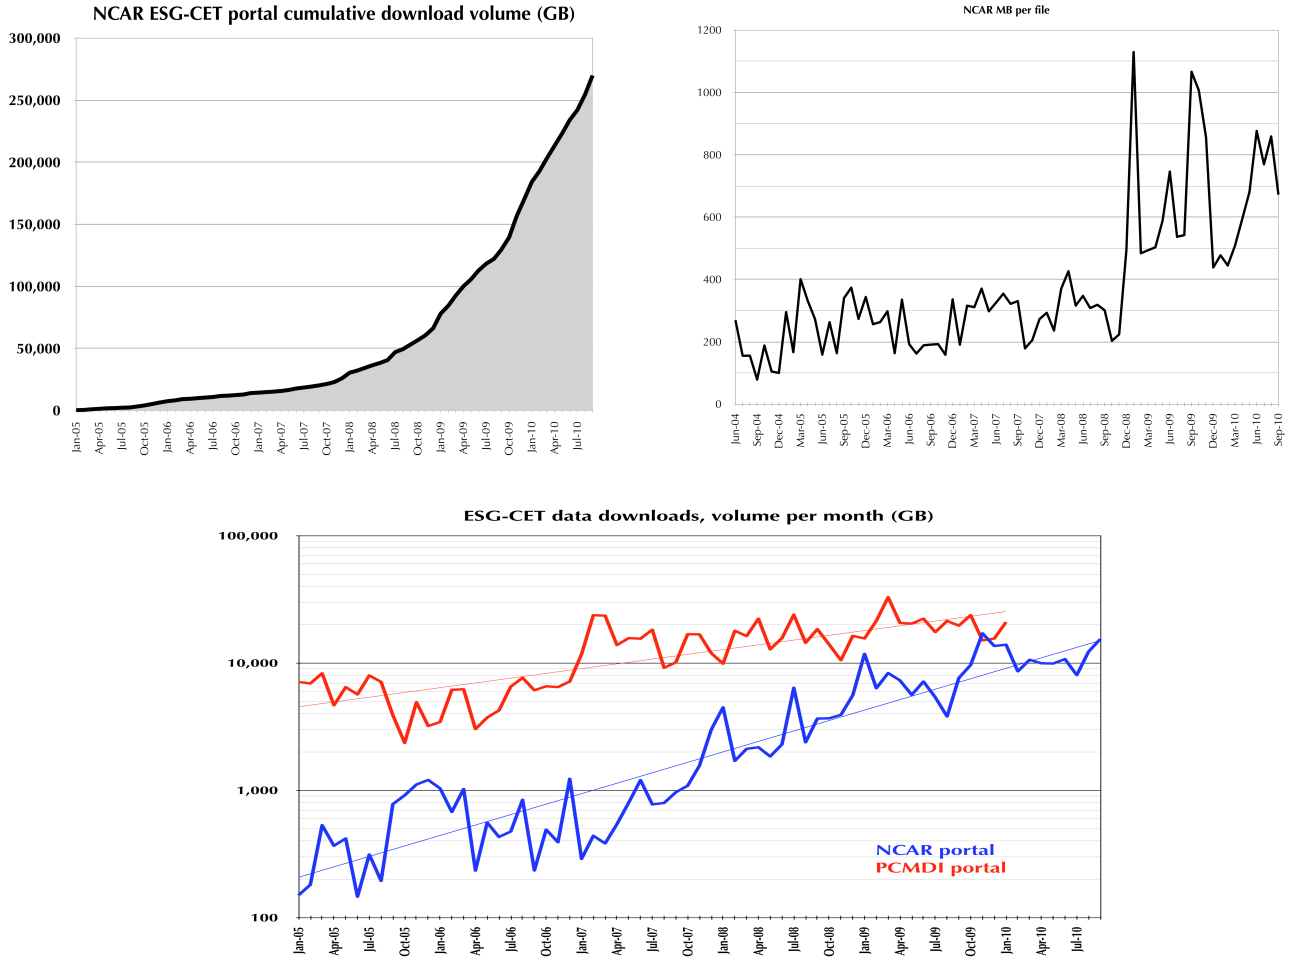


**Figure 7:** Additional download metrics.

## 3.2 Data Node Software Development

The ESG Data Node software team continued work on the software stack, primarily focusing on Data Node security. Additional ESG collaboration with external partners such as BADC and DKRZ has accelerated and enhanced several Data Note components, such as security, replication, and versioning. The Data Node comprises multiple software packages and tools (see Figure 8) that work with the Gateway to provide overall server-side federated functionality. The ESG software stack is being released as version 1.0.0.0. In this four-digit numbering system, the first number binds all the releases under a single Data Node value. As shown in Table 3, the Data Node software stack comprises many components from Web applications to CDAT and third-party components. Table 3 lists most of the applications formalizing version 1.0.0.0.

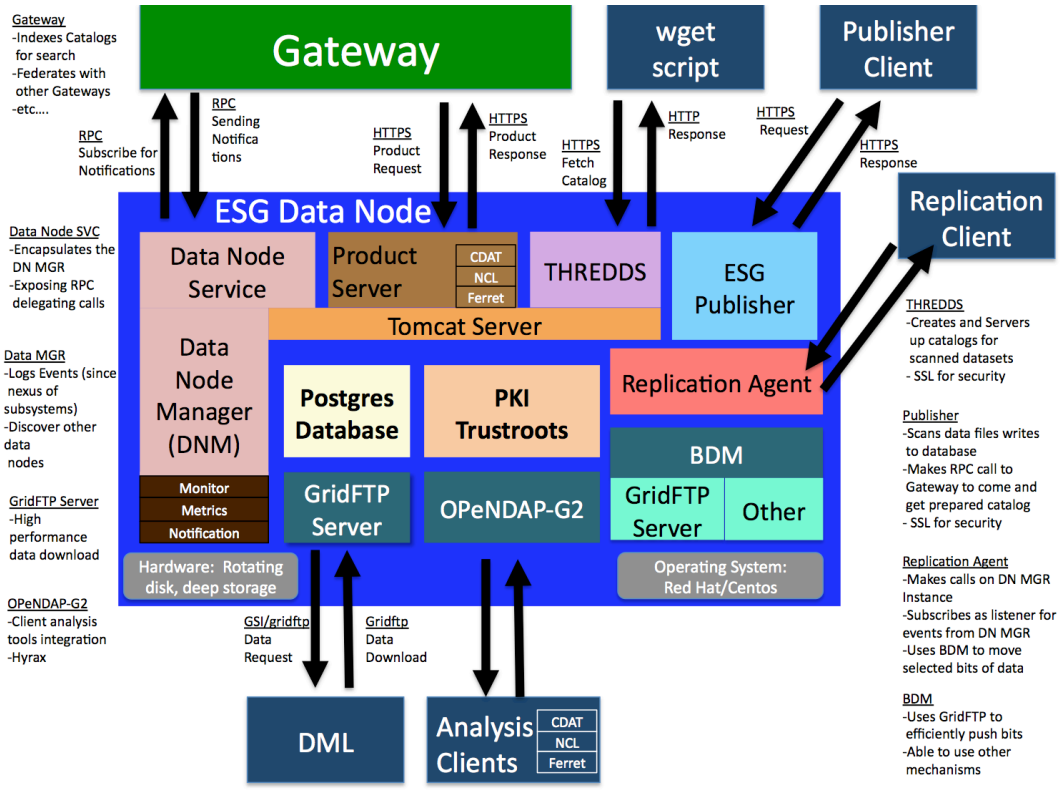


**Figure 8:** This diagram shows the diverse, complex operations performed by the Data Node within the ESG infrastructure.

**Table 2:** Software comprising the ESG Data Node software stack, version 1.0.4.0.

| Application | Name | Version Number |
|---|---|---|
| **Data Node (Web Application)** | ESG Data Node Software Stack | 1.0.0.4 |
| | | |
| **Language** | Java | 1.6.0.20 |
| **Language** | Python | 2.6.5.0 |
| | | |
| **Custom Component** | Node Manager | 1.0.0.4 |
| **Custom Component** | Publisher | 2.6.3.0 |
| **Custom Component** | Versioning (Publisher) | 1.0.0.0 |
| **Custom Component** | Replication (Publisher) | 1.0.0.0 |
| **Custom Component** | Bulk Data Mover (Replication) | 1.0.0.0 |
| **Custom Component** | Product Server (LAS) (**future release**) | 7.1.2.0 |
| | | |
| **Security** | SAML | 1.1.3.0 |
| **Security** | MyProxy | 4.9.0.0 |
| **Security** | OpenIdRelyingParty | 1.2.1.0 |
| **Security** | Globus | 4.2.1.0 |
| | | |
| **Analysis Tool** | Ferret | 6.6.4.0 |
| **Analysis Tool** | CDAT | 5.2.0.0 |
| **Analysis Tool** | NCL (**future release**) | 5.2.1 |
| | | |
| **3rd-Party Library** | HDF5 | 1.8.4-patch1 |
| **3rd-Party Library** | NetCDF | 4.1.1 |
| **3rd-Party Library** | Readline | 6.1.0.0 |
| **3rd-Party Library** | Xml2 | 2.7.6.0 |
| **3rd-Party Library** | Xslt | 1.1.26.0 |
| **3rd-Party Library** | Flex | 2.5.35.0 |
| **3rd-Party Library** | NumPy | 1.4.1.0 |
| **3rd-Party Library** | Pmw | 1.3.0.0 |
| **3rd-Party Library** | PyOpenGL | 3.0.1.0 |
| **3rd-Party Library** | Setuptools | 0.6c11 |
| **3rd-Party Library** | Bison | 2.4.0.0 |
| **3rd-Party Library** | Pam_Pgsql | 0.7.0.0 |
| | | |
| **3rd-Party Software** | Ant | 1.8.1.0 |
| **3rd-Party Software** | Git | 1.7.1.0 |
| **3rd-Party Software** | Curl | 7.20.1.0 |
| **3rd-Party Software** | Gsoap | 2.7.15.0 |
| **3rd-Party Software** | Postgres | 8.4.3.0 |
| **3rd-Party Software** | Tomcat | 6.0.26 |
| **3rd-Party Software** | THREDDS | 4.2.1.0 |
| **3rd-Party Software** | GridFTP | 3.24.0.0 |
| **3rd-Party Software** | Tcl/Tk | 8.5.8.0 |

Data Node development highlights for finalizing the delivery of CMIP5 (IPCC AR5) data by October 20, 2010, include the following:

- ***Simplified Data Node installation process:*** We codified the Data Node installation process into an interactive script that will fully install and configure all components. Written in BASH (a Unix command shell) for portability, the script also includes flags for security management and testing.
- ***New security using SAML and OpenID (tokenless):*** We revised the security infrastructure for tokenless authentication. The new framework uses public key encryption throughout to ensure communication and data integrity. In addition, the Security Assertion Markup Language (SAML) standard is employed for information exchange among federation members. The older token-based system is still supported for compatibility with older clients.
- ***Secure attribute service support using SAML with Yadis service:*** The attribute service leverages Yadis and uses SAML as the transport protocol. The attribute service provides secure query responses to the federated entities.
- ***DRS support for file system layout:*** DRS, a naming taxonomy for climate data files, allows users to construct canonical full-path names to data files in the system.
- ***Solr search support:*** The Data Node has a Solr instance that contains an index of the metadata gleaned from the scanned netCDF data files. The Solr index speeds up ingest and search operations over the metadata collection. Deploying this open-source search platform allows users to rapidly search the metadata.

For future releases, the Data Node team will work on:

- ***DRS file-system indirection layer for support of novel file systems:*** We are working to provide a layer of indirection between the naming of a file or set of files and the location of the data in a given data store. Using the DRS taxonomy as a file-system-like structure is a natural and canonical way of naming files so that the system can resolve a file name into that file's actual location in the back-end data store. This layer of indirection allows for sophisticated and novel file-system architectures, which may be more amenable to dealing with the magnitude of data and the distributed nature of the ESG federation.
- ***Virtual machine (VM) packaging of the Data Node:*** Packaging the entire Data Node as a VM provides several positive features to the system. First, the system configuration and installation issues are greatly reduced. Second, leveraging VM technology allows us to build more fluid systems and workflows that can take advantage of the cloud-computing model for work distribution and computation.
- ***Enhanced Data Node metrics interface:*** Each node will have a metrics interface served over the Hypertext Transfer Protocol (http) to show the state of that machine. A number of aggregating interfaces will provide administrators with a bird's-eye view of all the machines in the system, allowing them to drill down to specific data nodes and even files, as needed.
- ***Peer-to-peer node discovery and ad hoc overlay mesh network of nodes:*** Leveraging peer-to-peer discovery and ad hoc networking protocols and tools will further provide fluidity to the system. It also will allow the system to commission ad hoc groups of machines based on performance and/or storage criteria for performing different computational tasks. The peer-to-peer interconnection layer coupled with the fluidity provided from the VM layer creates a potent platform for performing data-intensive distributed computation at tremendous volume.

- ***Distributed index search across domains:*** The distributed nature of the ESG system requires that data be easily queried for and found across all the constituent member nodes. Given the magnitude of data involved, we will provide a distributed index such that data can be quickly found and fetched for further analysis or computation.
- ***Update to ESG Publisher's graphical user interface:*** Providing a stand-alone interface for publishing data to the system is a key requirement for ESG-CET. We are developing a publishing system independent of the nodes that will hold the data. Our intuitive, easy-to-use graphical user interface (Figure 9) will be distributed via JAVA Web Start, allowing scientists to publish experimental results from their home sites to the ESG federation. This portable, dynamic interface will allow researchers to scan and publish large volumes of data in a single operation, thus quickly making their climate data available worldwide.

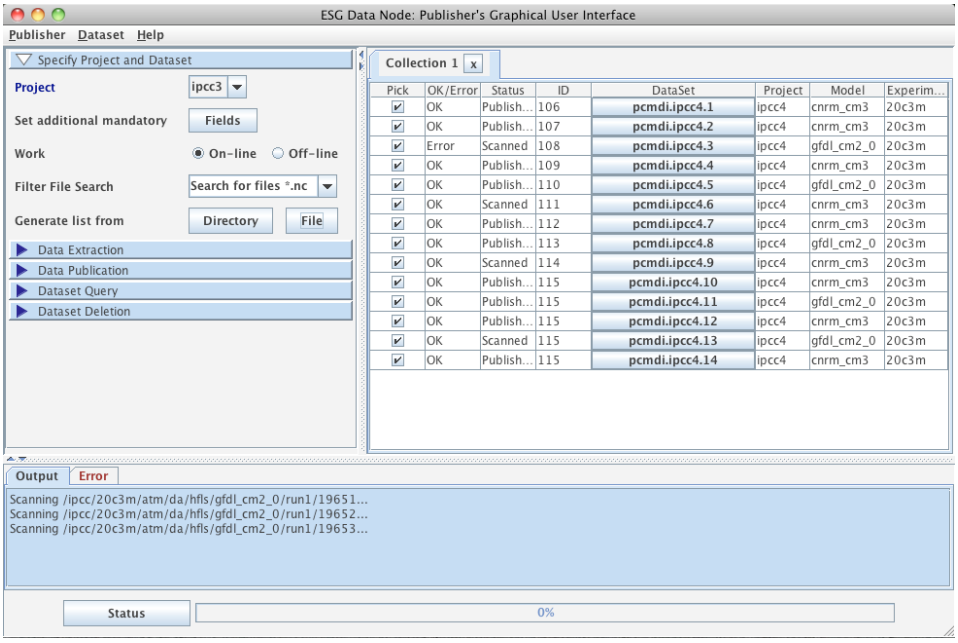


**Figure 9:** This graphical user interface is easy to navigate and allows users to scan and publish large volumes of data in a single operation.

- ***ESG Data Node Dashboard:*** The ESG Dashboard allows operators to monitor the Data Node in real time as data are added through the ESG Publisher. Data requests are sent from the client to the server with Asynchronous JavaScript and XML (AJAX). Data returned from the server is formatted in JavaScript Object Notation (JSON) and displayed in sortable tables and charts using a combination of JQuery and Hypertext Markup Language (HTML). Information on disk space, central processing unit (CPU) usage, memory, network resources, load, and peer connections are presented in a comprehensive and intuitive Web interface. As shown in Figure 10, pie charts monitor disk space, and line graphs chart CPU memory usage over time, while tables display user=interaction statistics.

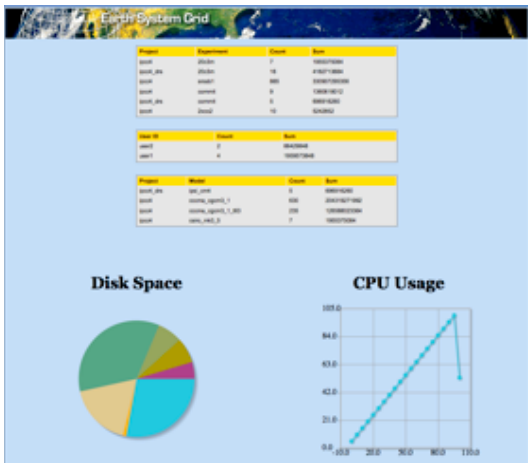


**Figure 10:** This simple interface is easy to navigate and allows data operators to have a unified view of critical data node resources and reports displayed for maximum efficiency.

## 3.3 ESG Server Setup and Testing

ORNL has deployed a Gateway and two Data Nodes on the ESG. One Data Node has access to approximately 150 TB of disk storage, and the second is designed to serve files from ORNL's High Performance Storage System (HPSS) tape archive (see Figure 11). The disk storage is provided by a Lustre file system consisting of three servers and an LSI XBB2 disk array. The servers are connected to the file system with DDR InfiniBand. The data nodes and gateway have 10-Gb and 1-Gb Ethernet connections to ESNet, respectively.

The first data node serves data from the LSI disk array, while the second data node is designed to serve data from HPSS. Because of ORNL's security restrictions, accessing HPSS requires two separate servers running customized software, as described below.

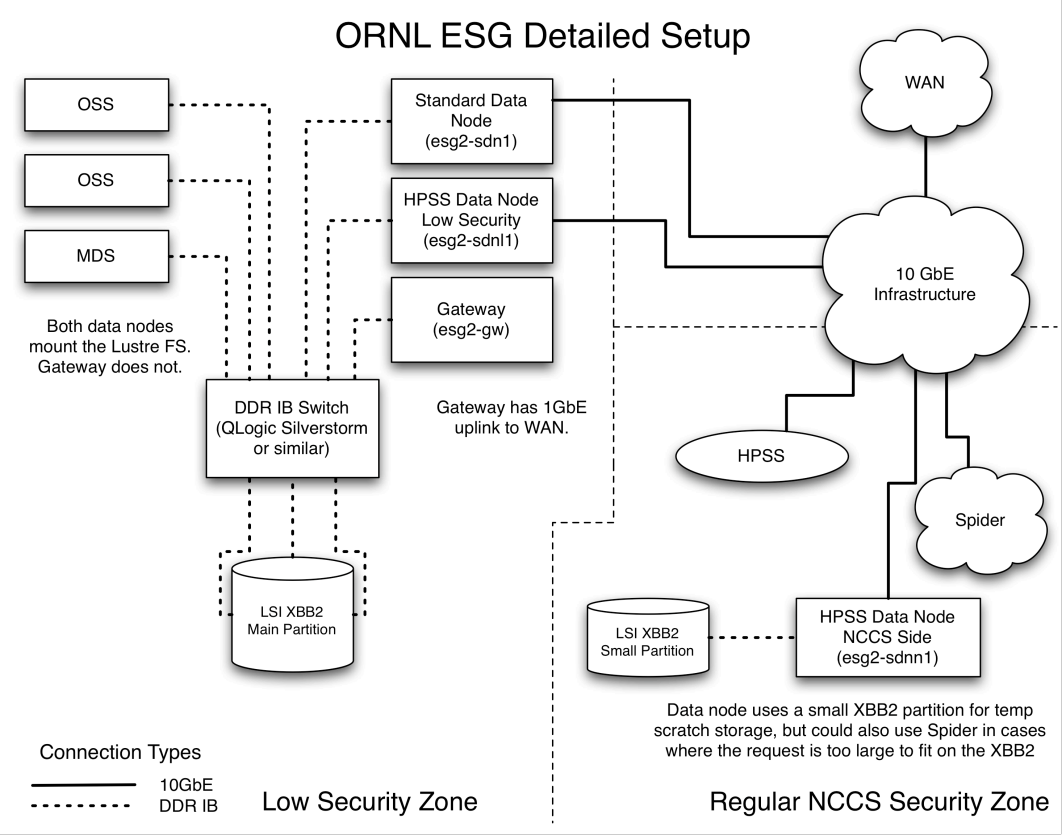


**Figure 11:** ORNL ESG setup.

The Gateway and the standard data node use the regular ESG software stack with very little in the way of customization. These systems were installed while LLNL was working on an automated installation script, so we were able to beta-test the script and help debug various issues during deployment.

Once the gateway and standard data nodes were operational, we coordinated with NCAR to test metadata exchange between our two Gateways. This test went well and uncovered no serious bugs. We successfully imported metadata from the NARCCAP project on the NCAR Gateway as well as C-LAMP data from the ORNL Gateway.

After testing the metadata exchange, we tested the OpenID federation, again between ORNL and NCAR, and the system performed as designed. The ORNL gateway will accept users who have been authenticated by the NCAR Gateway OpenID server, and the NCAR Gateway will accept users authenticated by the ORNL Gateway OpenID server.

## 3.4 ESG Deep Archive Challenges

The HPSS deployed at ORNL can store multiple petabytes of data for long-term archival purposes, while ESG's online disk capacity at ORNL is limited to hundreds of terabytes. Modeling activities at ORNL are projected to exceed the capacity of the disk cache and require the use of our HPSS archive. When deep archives are accessed from ESG, the system typically uses only the HPSS storage resource manager (SRM) within BeStMan; however, ORNL's multilevel security environment made this standard approach impossible.

ORNL's HPSS archive is in a medium-security enclave requiring two-factor authentication, while the ESG gateway is in a low-security zone that does not require two-factor authentication. Systems in ORNL's medium-security enclave are not allowed to have any ports open to the public Internet. Another challenge in this configuration is limiting data access so that the ESG HPSS low-security data node can access only the data published to ESG.

In addition to these challenges, enhancements to the ESG data publication process were required to provide seamless access to data published from ORNL's deep archive to ESG. Prior to making these improvements, data served through the HPSS SRM and published to ESG captured and displayed only the file name, size, and location along with other trivial attributes, whereas data published on a local disk included virtually all properties of the netCDF file (e.g., variable names and units). ORNL's enhancements to the publication process bridged this gap, allowing data published via the HPSS SRM to harvest all properties of the underlying data set on the ESG Gateway.

### *3.4.1 ESG HPSS Access: Design and Implementation*

The "standard" ESG method of integrating with SRM is to have a data node with direct access to HPSS (or other mass storage). The data node runs BeStMan (or some other SRM implementation), and the ESG Gateway knows the URLs of the files that BeStMan can serve. The user selects files from a list on the Gateway, which returns a list of URLs. The user then uses SRM client software to retrieve those URLs from the data node.

We cannot adapt the standard implementation, however, because doing so would expose all of ORNL's HPSS archive to additional attack vectors from the public Internet. We solved this problem by implementing a flexible client–server architecture: one low-security data node server (L-DN) residing within the low-security zone and another medium security data node server (M-DN) residing within the medium-security enclave (see Figure 12). Although the medium-security server can communicate with HPSS, it is isolated from the public Internet. The low-security server resides on the public Internet and provides a place for ESG clients to connect to and download files; it cannot however, communicate directly to ORNL's HPSS infrastructure.

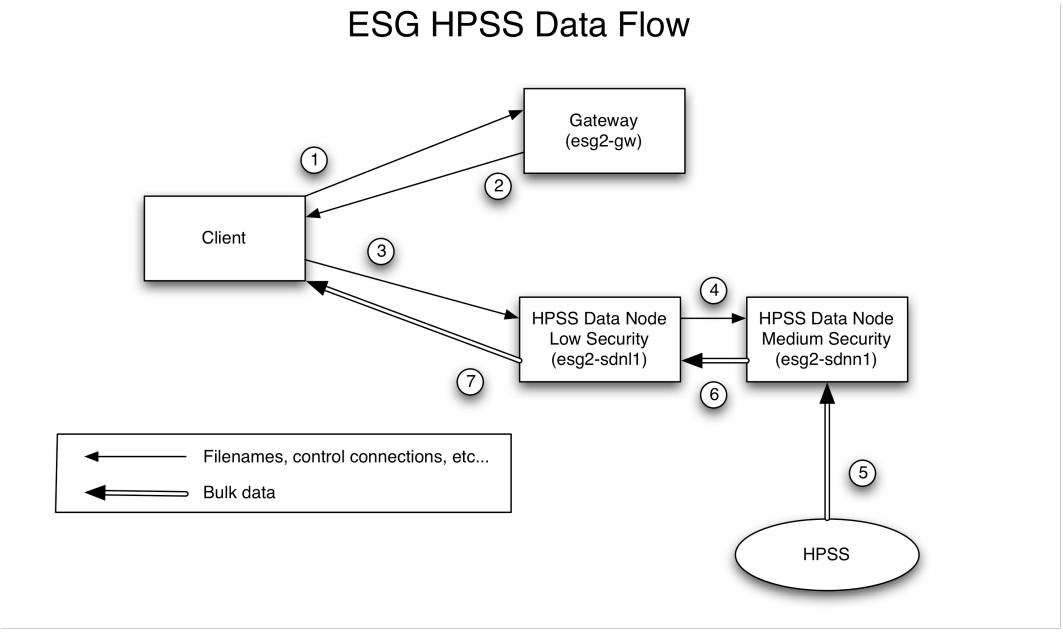


**Figure 12:** ESG HPSS data flow.

To enable this dual server design, we implemented the following components and customizations:

- New Listener program that can correctly parse the SRM output of NCCS.
- New plug-in module that leverages the data node stack extension capability.
- New BeStMan plug-in for handling SRM copy requests from the Gateway and interacting with the back-end server.
- New back-end server that runs as daemon process on the medium-security enclave, which handles SRM requests from BeStMan, interacts with HPSS server for local data transfer, and invokes GridFTP client for remote data transfer.
- Definitions for a set of XML-RPC-based API calls that serve as the communication protocol between the BeStMan plug-in and the back-end server.
- White-list support on the back-end server that verifies every request coming from the low-security zone against an authorized published file list.

With these new components in place, ORNL has developed an end-to-end infrastructure that provides access to data sets within ORNL's HPSS archive while providing rich metadata similar to that of data sets published from local disk. Figure 13 illustrated the data sets available on the ESG Gateway residing on ORNL's HPSS archive. These data sets were published using ORNL's enhanced data-publication mechanism. Figure 14 illustrates the rich metadata made available to the user via this publication mechanism. Metadata for individual data sets residing on the deep archive are no different from the metadata for data sets on the local disk.

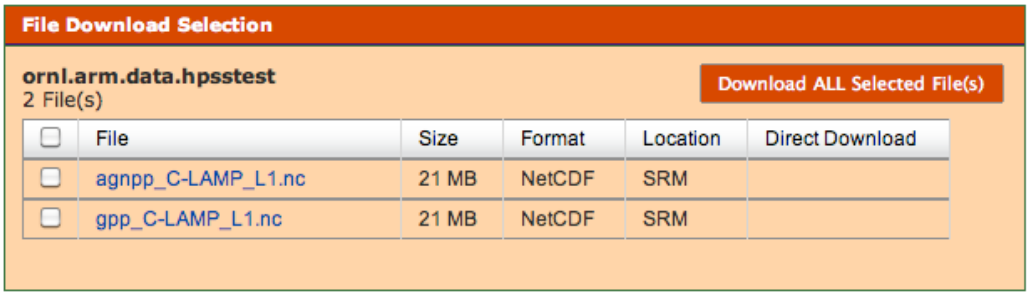


**Figure 13:** Examples of data sets to be published to ORNL's ESG Gateway from deep archive.

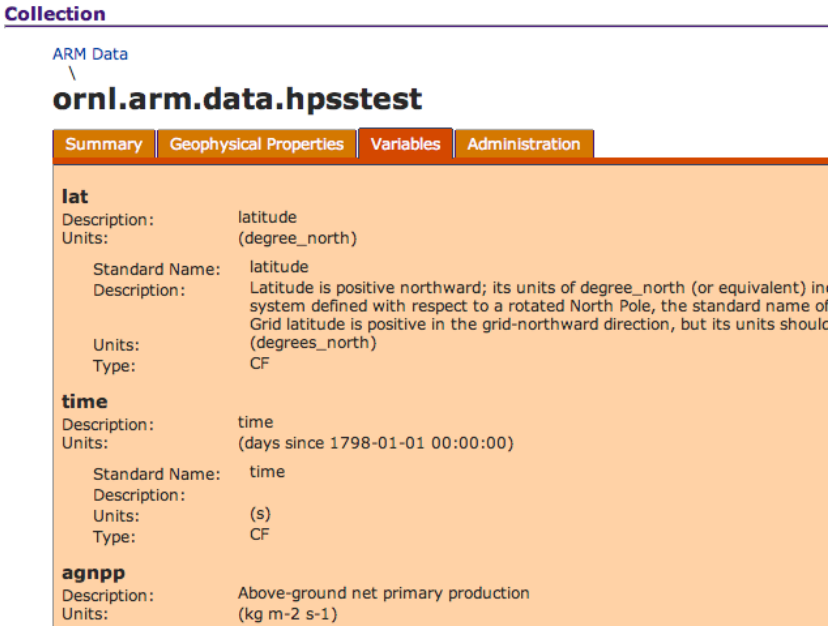


**Figure 14:** Full metadata support is provided for data sets published from deep archive including descriptions for each variable in the netCDF file.

When users request a data set for download, the system generates a request to L-DN. M-DN periodically queries L-DN, and when it discovers a new request, M-DN begins processing the download. As part of this operation, M-DN checks the request against a white list and then transfers the data set to the local disk cache. M-DN invokes a GridFTP transfer of the data set from local disk cache to L-DN and then notifies L-DN that the request has been completed. Once the data set is on L-DN, the user is notified that the set is available for direct download as illustrated in Figure 15.

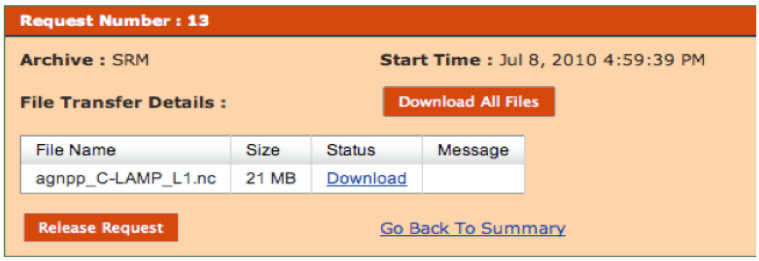


Figure 15: Dataset on deep-archive available for download via ORNL's BestMand plugin and end-to-end infrastructure

## 3.5 Cyber Security and Framework

The attribute framework in the ESG federation supports site and VO attributes. We worked with the ESG team to standardize formats for how these attributes are communicated within the federation. Argonne provided samples and framework code for generating attributes in the agreed format and embedding them in certificates for communication. Recent efforts to standardize the attribute names, such that there are agreed namespaces across the federation and attribute authorities associated with the attribute names, are designed for federation-wide exchange of attributes and attribute authority enforcement at the consumer level.

JPL and ANL in collaboration with BADC worked on supporting, upgrading, testing, and deploying the core packages of the ESG security infrastructure, including the SAML-based Attribute and Authorization Services, and the OpenID Relying Party. Specifically:

- The ESG security packages for SAML-based services and OpenID single sign-on (SSO) have been structured as modular, reusable libraries that can be deployed within an ESG Gateway, an ESG Data Node, or any other Java-based application.
- A framework we developed secures generic Data Node applications and resources by deploying and configuring non-invasive front-end application filters. The framework is currently used to secure the Thredds Data Server with ESG security, and progress is under way to apply it to LAS.
- The Java-based ESG security packages have been tested for interoperability with the Python-based installations deployed at BADC, based on the adoption of common protocols and APIs.

Ongoing work to secure LAS will eventually integrate the product server with the ESG-CET data node stack. This effort will allow product access using LAS to enforce the same permissions and policies that restrict download. Argonne played a key role in developing the architecture for this security solution and is working with other ESG-CET developers to prototype and implement these solutions.

## 3.6 Berkeley Storage Manager

BeStMan is the SRM implementation used in ESG. BeStMan's architecture is based on a modular design that makes it easy to interface to various storage systems, including a Mass Storage System (MSS). We developed versions of BeStMan for the HPSS at LBNL/NERSC and at ORNL and for the MSS at NCAR. Developed over the last 2 years, the three versions deal with different security mechanisms. During this past reporting period, we deployed BeStMan servers at LBNL/NERSC,

NCAR, and ORNL to serve ESG users in production through the new ESG Gateway portal. The ESG Metadata Publishing tool deployed at these three sites also uses BeStMan servers.

In addition, we also made the following feature additions and bug fixes during this reporting period:

- Added support for setting the maximum log file sizes.
- Improved efficiency of the parsing mechanism for the NCAR MSS plug-in.
- Fixed the MSS request-processing queue to resolve hanging issue while accessing the queue.
- Updated the MSS plug-in parsing mechanism to support additional use cases.
- Fixed the request monitoring thread for long-running requests.

## 3.7 DataMover-Lite

Changes in requirements and suggestions from user experiences for DataMover-Lite (DML) came up during this period for its operation with the authorization plugin in gridftp servers in the new ESG federated system. The main requirement is to use the common ESG directory for user proxy and certificates so that duplicate efforts among ESG components can be minimized. DML version 3 is now fully integrated with the ESG authorization model and plugins. The stand-alone version and the web start version have been used in production. The following tasks were accomplished:

- Reset button has been added for resetting the configuration parameters.
- Support for common ESG configuration directory has been added for user proxy file and bootstrapped certificates from MyProxy server.

## 3.8 Bulk Data Mover

- We released BDM and are testing it extensively with the ESG Replication Service. During this reporting period, we also added the following features to improve BDM efficiency.
- Support for parallel browsing, a feature that accelerates the source directory browsing for collecting files information.
- Support for passive connections and run-time option in file browsing and file transfers with firewall setup.
- An optional feature for transfer performance testing.
- Support for destination space checking for limited storage.
- Capability to reuse network connections for data transfers from remote source directory browsing.
- Support for root paths and absolute and relative paths in URLs.
- Capability to add and remove network connections dynamically during active transfer requests for adaptive data transfers, based on the bandwidth availability and network performance.

In addition, we tuned the back-end database for improved performance in storing information on requested files and the status of those operations, and we improved the screen output BDM operations to provide more concise, user-oriented information. We also implemented a more efficient approach for processing existing destination files when a request to resume is submitted or when the system needs to

synchronize file information between the source and destination. This improvement was accomplished with parallel-multithreaded checks that depend on the overwriting or non-overwriting privileges established for the destination files.

## 3.9 Data Replication

The USC/ISI team successfully implemented the ESG replication client during this reporting period and performed end-to-end testing of the client to mirror data sets from LLNL to sites at USC/ISI and NCAR.

The replication client typically runs on a data node that will store a replica of an existing data set (the mirror data node). The replication client is given the name of a data set to be replicated and begins by querying the metadata catalog on a gateway to determine the list of files in that data set. The replication client also queries the gateway for a pointer to the THREDDS catalog on the data node where the data were originally published (the source data node). Next, the replication client prepares data transfer requests to copy data from the source node to the mirror node. This data transfer can be performed using either the BDM client or the GridFTP globus-url-copy client. The replication client then copies the THREDDS catalog from the source data node to the mirror. After successfully copying the data set, the replication client invokes the publication API developed at LLNL to scan the copied files and create a THREDDS catalog for the replicated set. Next, the replication client software compares the newly generated THREDDS catalog for the replicated data set with the THREDDS catalog from the source data node and verifies that essential elements are consistent in both catalogs. If this check succeeds, the client then invokes the publication API to publish the replicated data set to a gateway, where it can be discovered and accessed by ESG users.

Specific accomplishments of this period include the following:

- Implemented, tested, and documented the functionality of the replication client that queries the gateway metadata server and prepares a control file (XML for BDM) or an executable shell script (for globus-url-copy) to download files for a requested data set. Created a simple database to manage the originating data node THREDDS catalog so that it could be used later for comparison with the newly generated THREDDS catalog for the replicated data set.
- Implemented, tested, and documented the functionality that invokes the ESG-CET publication API developed at LLNL to scan downloaded files, create a THREDDS catalog, and if the catalog comparison is successful, publish the replicated data set.
- Designed, implemented, tested, and documented code to compare two THREDDS catalogs: the catalog from the source data node and the newly generated catalog from the mirror data node. A candidate replica data set will not be published to the gateway unless the comparison succeeds, indicating that the required information is present in the THREDDS catalog for the replicated data set.
- Packaged the replication client and documented its dependencies for easy installation at NCAR.
- Ran end-to-end replication tests on data sets ranging from less than 1 up to 50 GB in size. These tests mirrored data sets from LLNL to a data node at USC/ISI. The USC/ISI team also worked closely with NCAR personnel to help install the replication client at NCAR. The two teams tested different portions of the replication logic and performed end-to-end replication of the test data sets using the globus-url-copy data transfer client.

Plans for the coming 6-month period are to complete testing at NCAR using BDM for data transfer; provide support to our European partners to deploy and test the replication client functionality; hold follow-on discussions to determine additional requirements for replication; and implement these as appropriate for the ESG and GO-ESSP collaborations.

## 3.10 Product Services

Naturally, many of the data fields housed in the ESG-CET data archives are vectors. Vector components are typically stored in separate files and do not receive special treatment when they are stored in netCDF. Furthermore, the CF conventions provide little in the way of support for identifying vectors among a collection of data files. During the reporting period, the PMEL LAS team developed and implemented a set of heuristics that allow the software collecting configuration metadata to automatically recognize and configure vectors into LAS. LAS can then provide products for plots and analysis of the individual components, and a user can see and select new vector variables from the list of available data. When a vector variable is selected, the user sees a new collection of appropriate visualization tools, including the comparison interface vizGal, shown in Figure 16. This interface allows users to make a gallery of vector plots. For example, it can compute the differencing of two vector fields.

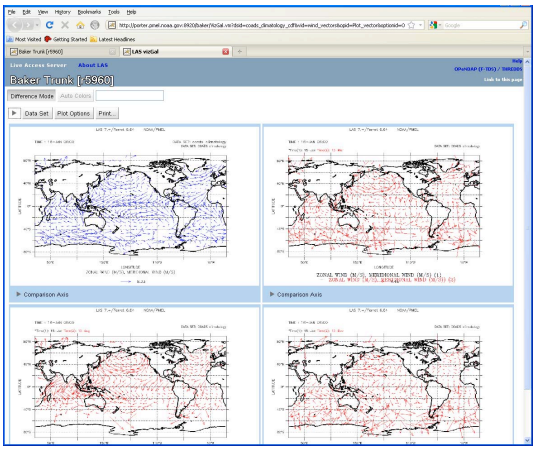


**Figure 16:** A browser displaying the vizGal interface with a vector comparison plot

We also enhanced the interface and back-end infrastructure of LAS so users can select two or more variables. Also, when more than one variable is selected, the user can work with a special collection of visualization tools appropriate for the number of variables selected. For time-series plots, selecting more than one variable creates a single plot of the selected variables with appropriate scaling and multiple vertical axis labeling and tick marks. A user can make property–property scatter plots of two data variables, and the Ferret engine will automatically compute and display fit lines and statistics in property-property plots. The user can also choose to color scatter markers according to the value of a third variable.

The LAS team in collaboration with teams at NASA and ORNL has successfully applied the authentication infrastructure used by ESG-CET to LAS installations. Users seeking to make plots from

an ESG-CET LAS are automatically redirected through the appropriate OpenID authentication. The team also applied the authentication filters used for ESG TDS installations to LAS installations deployed as confluence servers on a Gateway node and as regular LAS servers on data nodes. Work is still in progress on aspects that will restrict access to data and products for which the user has authorization. In addition, the LAS team updated LAS so that vector plotting and comparisons are automatically enabled. With this functionality, users can create multivariable time-series, property–property and property–property–property plots in LAS (see Figures 17–19). More detailed information on this functionality is available at the Web site for the LAS project: http://ferret.noaa.gov/LAS.

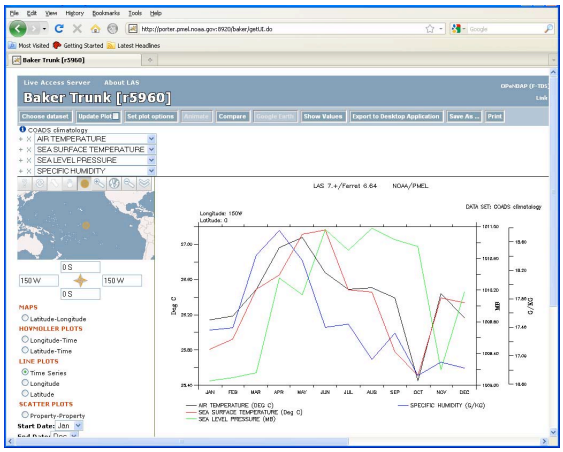


**Figure 17:** This screen shot shows a time series of multiple variables plotted in LAS with multiple vertical scale information displayed and documented.

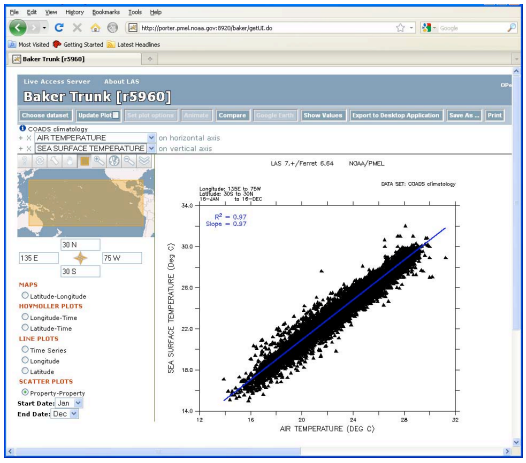


**Figure 18:** A Web browser display shows an LAS property–property plot.

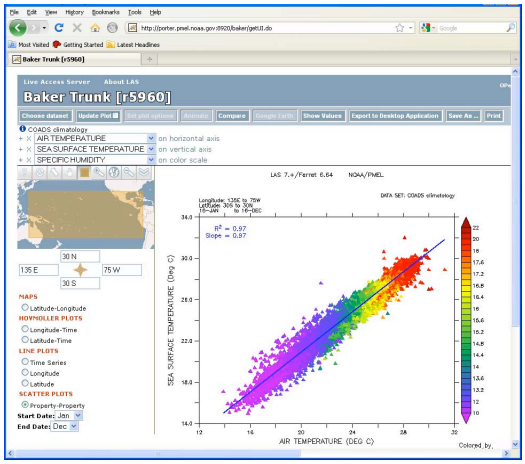


**Figure 19:** This Web browser display shows an LAS property–property plot colored by a third variable.

Ongoing work to secure the LAS will result in the integration of the product server within the ESG Data Node software stack. This effort will allow product access using LAS to enforce the same permissions and policies that restrict downloads. ANL is playing a key role in architecting the LAS security solution, and are working with other developers on the ESG-CET team to prototype and implement these solutions.

# 4 ESG-CET Group Meetings

The ESG-CET executive committee holds weekly conference calls on Mondays at 9:00 a.m. Pacific time. At these meetings, the committee discusses the priorities and issues that will make up the agenda for the weekly project meetings, which are held on Thursdays at 12:00 p.m. Pacific time via the Access Grid (AG). At the AG meetings, the entire team discusses project goals, design and development issues, technologies, timelines, and milestones. Given the need for more in-depth conversation and examination of work requirements, several face-to-face meetings are planned for the next reporting period.

# 5 Collaborations

To effectively build an infrastructure that can accommodate petascale data management and analysis, we established connections with other DOE Office of Science SciDAC-funded projects and programs at various meetings and workshops, such as the SciDAC 2010 Conference held in Chattanooga, Tennessee. In particular, we established collaborations with the following groups: TeraGrid Science Gateways, Earth System Curator, NOAA's GIP, METAFOR, World Meteorological Organization (WMO) Information System, Scientific Computing and Imaging Institute at the University of Utah, SciDAC VACET, SciDAC SDM, SciDAC CEDPS, Southern California Earthquake Center, Tech-X Corp., NASA Langley, NASA Goddard, GO-ESSP, and many others.

## 5.1 2009 Global Organization for Earth System Science Portal Workshop

Dean N. Williams, Steve Hankin, and Don Middleton are members of the GO-ESSP steering committee responsible for organizing the 2011 workshop, which will be held in Asheville, North Carolina, USA, and hosted by NCDC.

The GO-ESSP workshop facilitates the organization and implementation of an infrastructure for full data sharing among a consortium spanning continents, countries, and intergovernmental agencies. This consortium envisions an environment that allows users open access to petabytes of model-generated, satellite, and in situ data including physical, biogeochemical, and ecosystem content. All initial ESG Federation partners attended the 2010 workshop (i.e., LLNL, NCAR, GFDL, BADC, DKRZ, and the University of Tokyo).

## 5.2 NOAA Global Interoperability Program

NOAA's GIP program promotes coordination of software infrastructure development across agencies, the weather and climate communities, modeling and data services, and research and operational centers.

The Curator project, now hosted under GIP, continues to be a highly productive collaboration with ESG-CET. Curator serves as an active liaison between the EU-based METAFOR project, which is developing a CIM for CMIP5/IPCC, and ESG-CET, which will be delivering data to the global community. We have jointly demonstrated new "model trackback" capabilities, which effectively combine the METAFOR questionnaire with the query and browse functions of ESG.

## 5.3 WMO Information System

The WMO Information System (WIS) Program is developing a next-generation, globally federated data system to serve all of the WMO areas (such as climate, weather, and hydrology). Don Middleton contributes to several WIS committees and regularly provides briefings that encompass ESG-CET and our progress toward CMIP5/IPCC objectives and climate research in general. Our Gateway technology will provide federation with WIS while also serving as a WIS validation platform.

## 5.4 NSF TeraGrid Science Gateways Program

ESG-CET, NOAA GIP (Curator), and Purdue University staff collaborated on a proposal to develop a prototype Environmental Science Gateway and was granted support from TeraGrid's Science Gateway program for two years. The basic objectives of this project include developing a Gateway-based model execution capability, a supporting workflow system that can operate on the TeraGrid Gateway, and data publishing and harvesting capabilities that can interoperate with an integrated rule-oriented data system–(iRODS-) based archive as well as disk-hosted model output. Work on this effort continued throughout this reporting period.

# 6 Outreach, Papers, Presentations, and Posters

This section describes the outreach activities for this reporting period in the form of papers published and talks and posters presented.

## 6.1 Papers

### *6.1.1 BER Network Requirement Workshop*

Dean N. Williams, Gary Strand, and Galen Shipman, "Climate Science: Enabling Worldwide Access to Petascale Climate Data," *Proceedings of the Office of Biological and Environmental Research Network Requirements Workshop*, April 29–30, 2010. According to the report, the biggest BER user of ESnet in the future will likely be the climate community. Through ESG, LLNL and ORNL will host petabyte archives that will be accessible worldwide.

### *6.1.2 SciDAC '10 Conference Proceedings*

D. N. Williams, J. Ahrens, R. Ananthakrishnan, M. Balman, G. Bell, S. Bharathi, D. Brown, M. Chen, A. L. Chervenak, L. Cinquini, R. Drach, I. T. Foster, P. Fox, S. Hankin, D. Harper, N. Hook, P. Jones, D. E. Middleton, N. Miller, E. Nienhouse, R. Schweitzer, G. Shipman, A. Shoshani, F. Siebenlist, A. Sim, W. G. Strand, F. Wang, C. Ward, P. West, H. Wilcox, N. Wilhelmi, and S. Zednik, "Earth System Grid Center for Enabling Technologies: Building a Global Infrastructure for Climate Change Research," *Journal of Physics: Conference Series, SciDAC '10 Conference Proceedings*.

### *6.1.3 Proceedings of the International Symposium on Grid Computing*

A. Sim, D. Gunter, V. Natarajan, A. Shoshani, D. Williams, J. Long, J. Hick, J. Lee, and E. Dart, "Efficient Bulk Data Replication for the Earth System Grid," *Proceedings of the International Symposium on Grid Computing, Data Driven e-Science: Use Cases and Successful Applications of Distributed Computing Infrastructures* (ISGC 2010).

## 6.2 Talks

### *6.2.1 National Climate Assessment*

Dean N. Williams represented ESG-CET at the National Climate Assessment meeting sponsored by NSF, NOAA, and the U.S. Geological Survey (USGS), which was held at the USGS National Center Auditorium on September 21, 2010. As part of a more sustainable and ongoing effort, this group is looking to tie into the ESG infrastructure. A National Climate Assessment is required every four years, and the next report due in June 2013.

### *6.2.2 Fall 2010 Climate Visualization 10-05 Architecture Meeting*

Bob Drach and Dean N. Williams represented ESG at the Fall 2010 Climate Visualization 10-05 Architecture Meeting. This meeting brings together technologists and scientists interested in leveraging emerging technologies and solutions to support, enable, and enhance climate research. Meeting attendees were particularly interested in the ESG work on infrastructures and tools to support analysis and sharing of climate model output and observations.

### *6.2.3 Greenhouse-GAS Information Systems*

Dean N. Williams represented ESG at the Greenhouse-GAS Information Systems Requirement Meetings held at ORNL and Sandia National Laboratories in August and September. Further discussions are planned for FY11 regarding the potential role of ESG within the proposed mission operation center.

### *6.2.4 2010 Federation of Earth Science Information Partner (ESIP) Meeting*

Luca Cinquini, Dan Crichton, Cecelia DeLuca, Dean Williams presented, "The Earth System Grid Federation" at the 2010 ESIP Federation Meeting held at the University of Tennessee's Conference Center in Knoxville, TN July 20-23. The Federation of Earth Science Information Partners is a unique consortium of 121 organizations that collect, interpret and develop applications for remotely sensed Earth observation information. Luca represented the ESG team at the meeting.

### *6.2.5 Climate Science for a Sustainable Energy Future*

Dean N. Williams and Galen Shipman represented ESG at the Climate Science for a Sustainable Energy Future (CSSEF) Project meeting held at ORNL on May 24–26. Their talk, "Federation and Integration of Data from Multiple Sources," covered critical data functionalities, interagency partnerships, and requirements for CSSEF.

### *6.2.6 DOE ASCR–BER Network Requirements Workshop*

Dean N. Williams, Galen Shipman, and Gary Strand discussed OBER-related climate network needs at the ASCR–BER Network Requirements Workshop. Dean presented "Enabling Worldwide Access to Petascale Climate Data" and wrote a 16-page paper based on this talk for the workshop report. The workshop report discusses outstanding network issues, future needs, and scientific directions for the climate community. The planned network will have ESG as the cornerstone for distributing climate data to the community. In addition, the report notes that the climate research community is likely to be the biggest BER user of ESnet in the future. Climate sites (e.g., LLNL, ORNL) will host huge data sets that will be accessed by users worldwide. By 2020, the climate data repositories at these sites may exceed hundreds of exabytes in size. Several key sites for distributing these data are the CMIP5 Data Centers at LLNL, BADC, DKRZ, the Australian National University, and the Japan Agency for Marine-Earth

Science and Technology. NCAR and LBNL will distribute data from the CCSM model, and ORNL will distribute observational and other data. Data will be replicated across these and other locations, and bandwidth requirements will likely reach 100 Gbps in the next few years.

#### 6.2.7 DOE SciDAC UltraViz Institute All-Hands Meeting

Don Middleton gave an invited talk at the 2010 DOE SciDAC UltraViz Institute All-Hands Meeting, held at the University of California at Davis in April 2010. His talk, "Challenges in Climate Data Analysis and Visualization," included ESG-CET as a primary theme.

#### 6.2.8 DOE BER Climate Change Modeling Program Science Team Meeting

Dean N. Williams presented "The Earth System Grid Center for Enabling Technologies (ESG-CET): Serving Climate Data to the World" at the DOE Climate Change Modeling Program Science Team Meeting in Gaithersburg, Maryland, March 29–April 2, 2010.

#### 6.2.9 HPDC Challenges for Large Applications in Distributed Environments

Don Middleton presented "The Earth System Grid Federation: A Globally Distributed Environment for Climate Research," an invited talk, at the 2010 Conference for the High Performance Distributed Computing, Challenges of Large Applications in Distributed Environments, which was held in Chicago, Illinois.

#### 6.2.10 ESG Federation Meeting

Eric Nienhouse presented "Enabling Community Driven Development in the ESG Gateway," at the international ESG Federation meeting held in Princeton, New Jersey.

Gavin Bell presented "Peer-2-Peer Architecture" at the international ESG Federation meeting held in Princeton, New Jersey.

Bob Drach presented "CMIP5 (IPCC AR5) System Requirements" ant the international ESG Federation meeting held in Princeton, New Jersey.

Luca Cinquini presented "Solr and Federation Architecture" and the international ESG Federation meeting held in Princeton, New Jersey.

### 6.3 Posters

#### 6.3.1 DOE SciDAC Annual PI Meeting

Don Middleton, Dean N. Williams, and Ian Foster presented the poster, "The Earth System Grid Center for Enabling Technologies: A Petascale Data Service for the Global Community," at the annual SciDAC PI Meeting, held in Chattanooga, Tennessee, in July 2010.